EXHIBIT
E

UNITED STATES DISTRICT COURT

DISTRICT OF DELAWARE

SRI INTERNATIONAL, INC.,
a California corporation

Plaintiff and
Counterclaim-Defendant,

vs.

INTERNET SECURITY SYSTEMS, INC., a Delaware corporation; INTERNET SECURITY SYSTEMS, INC., a Georgia corporation; and SYMANTEC CORPORATION, a Delaware corporation,

Defendants and
Counterclaim-Plaintiffs.

_____/

CERTIFIED COPY

Case No. 04-1199 (SLR)

DEPOSITION OF GEORGE KESIDIS
VOLUME I

DATE: May 25, 2006

TIME: 9:13 a.m.

LOCATION: DAY CASEBEER MADRID & BATCHELDER
20300 Stevens Creek Boulevard
Suite 400
Cupertino, CA 95014

REPORTED BY: KAREN L. BUCHANAN
CSR No. 10772

8696
21416

Bell & Myers
CERTIFIED SHORTHAND REPORTER, INC.

50 AIRPORT PARKWAY, SUITE 205, SAN JOSE, CALIFORNIA 95110, TELEPHONE (408) 287-7500, FAX (408) 294-1211

haven't considered the question, so the answer would be I just don't know. 14:28:11

BY MS. MOEHLMAN: 14:28:12

Q. What is your understanding of how the impact analysis works in Fusion 2.0 and above? 14:28:16

A. I mean I -- I could give a very high-level description of what the module does at this point, but I would rather refresh my memory and just look at a reference, a high-level document describing it. I haven't -- we didn't really focus on that in our analysis, on that particular module, and so as a result, I'm just really not prepared to answer your question. 14:28:32 14:28:37 14:28:39 14:28:46 14:28:55 14:28:58 14:29:03 14:29:06 14:29:06

Q. Do you understand that SiteProtector Security Fusion module that existed before Version 2.0 consisted solely of the vulnerability impact analysis? 14:29:09 14:29:14 14:29:19 14:29:19

A. Yes, I do know that, yeah. 14:29:21

Q. And it is your opinion that those versions, in connection with the agents and SiteProtector, do not infringe, correct? 14:29:24 14:29:28 14:29:30

A. By inference, yes. By inference. 14:29:33

Q. So what, to your knowledge, changed between in the Fusion module 2 from the earlier version? 14:29:38 14:29:42

MR. POLLACK: Objection. Lacks foundation,

vague and ambiguous. 14:29:43

BY MS. MOEHLMAN: 14:29:46

Q. Do you know what's changed? 14:29:49

A. It was the conclusion of the correlation functions. 14:29:53 14:29:54

Q. It was the attack pattern core? 14:29:58

A. Right. 14:29:58

Q. Is that the element that you believe is relevant to the claims to the '615 and the '203 claims? 14:30:02 14:30:05 14:30:05

MR. POLLACK: Objection. Vague and ambiguous. 14:30:06 14:30:11

THE WITNESS: It essentially is why we -- we are contending security Fusion module 2.0. 14:30:15 14:30:15

BY MS. MOEHLMAN: 14:30:23

Q. Is it your understanding that the vulnerability impact analysis operates separately from the attack pattern component of Fusion? 14:30:28 14:30:34 14:30:40

MR. POLLACK: Objection. Vague and ambiguous. 14:30:40 14:30:43

THE WITNESS: My understanding is in principle, it need not. In principle, vulnerability assessments are used for detection. As a general rule, they can be used for detection, so -- even with respect to corroborating an attack pattern that you're 14:30:46 14:30:55 14:30:58 14:31:05

substantial evidence of generic code used -- or code between the agents and the attack pattern. 15:32:46 15:32:57

Q. If the Symantec construction -- I didn't mean to cut you off. 15:33:01 15:33:02

A. No. Go ahead. 15:33:04

Q. So I take it if the ISS constructions of "network monitor," "monitor," and "hierarchical monitor" were adopted, then there would be no infringement, correct? 15:33:10 15:33:14 15:33:20 15:33:21

MR. POLLACK: Objection. Vague and ambiguous. 15:33:26 15:33:26

THE WITNESS: No infringement among all the claims? For all the patents in suit? I'm not sure I -- again, we kind of tackled this on a claim-by-claim basis, but I would say a substantial number of the hierarchical claims, quote, unquote, hierarchical claims, again I found no evidence that SiteProtector or with Security Fusion 2.0 with the attack pattern module, that it used the generic code from the agents. 15:33:28 15:33:33 15:33:41 15:33:44 15:33:48 15:33:56 15:34:00 15:34:05 15:34:09

So with regard to Symantec's construction, the term is "dynamically configured." So you're asking is the attack pattern classifier dynamically configured and dynamically configurable. Is that -- I'm boiling it down to the right point you're making? 15:34:31 15:34:49 15:34:55 15:35:01

UNITED STATES DISTRICT COURT

DISTRICT OF DELAWARE

SRI INTERNATIONAL, INC.,
a California corporation,

Plaintiff and
Counterclaim-Defendant,

vs.

INTERNET SECURITY SYSTEMS, INC.,
a Delaware corporation; INTERNET
SECURITY SYSTEMS, INC., a Georgia
corporation; and SYMANTEC
CORPORATION, a Delaware corporation,

Defendants and
Counterclaim-Plaintiffs.
_____/

CERTIFIED COPY

CASE NO: 04-1199 (SLR)

DEPOSITION OF GEORGE KESIDIS
VOLUME II

DATE:	Friday, May 26, 2006
TIME:	9:00 A.M.
LOCATION:	DAY, CASEBEER, MADRID & BATCHELDER 20300 Stevens Creek Boulevard Suite 400 Cupertino, CA 95014
REPORTER:	Patricia Hope Sales, CRR CSR License Number C-4423

8705
21418

Bell & Myers
CERTIFIED SHORTHAND REPORTER, INC.

reports of suspicious activity." 11:18:25

Do you agree with SRI's construction of that 11:18:26

term? 11:18:30

A. Yes. 11:18:30

Q. What do you understand "functional unit" to be? 11:18:38

A. A -- I -- my opinion of a functional unit is -- 11:18:51

is similar to -- an instance of a functional unit is a 11:19:02

report of suspicious activity that -- as you would 11:19:04

receive from a network service monitor. 11:19:07

So that in turn, for example, you could 11:19:12

dispatch an integrated report or a meta-alert, say, to 11:19:18

a peer hierarchical monitor or a hierarchical monitor 11:19:33

at a -- at a higher layer. 11:19:41

So I would describe "functional unit" as, by 11:19:49

example, a meta-alert or a meta report of suspicious 11:19:52

activity. 11:19:59

Q. If instead of generating a meta report, data 11:20:03

that the -- data indicating -- let me start over. 11:20:11

If a hierarchical monitor, instead of 11:20:35

generating a report, displayed on a screen a group of 11:20:40

related reports, would that be in your opinion a 11:20:49

functional unit? 11:20:52

MR. POLLACK: Objection. Vague and ambiguous, 11:20:54

incomplete hypothetical. 11:21:00

THE WITNESS: I -- I would -- I would say no. 11:21:02

BY MS. MOEHLMAN: 11:53:36

Q. Okay. Did the console, the RealSecure console, in your opinion automatically receive the reports? 11:53:36 11:53:41

MR. POLLACK: Objection. Vague and ambiguous. 11:53:45

MS. MOEHLMAN: And it's "console" (pronunciation). I'm sorry. 11:53:47 11:53:49

THE WITNESS: Yes, I would agree. 11:53:52

BY MS. MOEHLMAN: 11:53:54

Q. Now, let's go to the claim construction statement. 11:54:00 11:54:03

If the ISS construction of automatically receiving and integrating the reports of suspicious activity were adopted, would the prior art RealSecure system in your opinion infringe that claim element -- or I'm sorry -- would it have that claim element? 11:54:10 11:54:16 11:54:21 11:54:27 11:54:34

MR. POLLACK: Objection. Vague and ambiguous, calls for speculation. 11:54:39 11:54:39

THE WITNESS: (Reviewing document(s).) 11:54:47

So I didn't opine on the ISS construction? 11:54:47

You are referring to the bottom of page three in -- 11:54:56 11:54:58

BY MS. MOEHLMAN: 11:55:00

Q. Yes, I am. 11:55:00

A. -- Exhibit 8, I believe? 11:55:01

Q. Yes. 11:55:03

A. The automatically receiving it? 11:55:03

The -- the kind of combination conducted by ISS, that is to say, merely displaying the events at a same console, is -- is not in my opinion what was meant by "integration" in the claim. 11:55:21 11:55:24 11:55:34 11:55:51

So I -- I'm assuming that if simply displaying the events as received is construed to be integrating, then I would agree that the -- the "automatically" element would be -- would be met, but I -- I didn't really -- haven't really thought about it too carefully. 11:56:08 11:56:17 11:56:29 11:56:36 11:56:41 11:56:43

Q. Is it your opinion that the RealSecure console in the prior art merely displayed the events as received? 11:56:45 11:56:51 11:56:55

MR. POLLACK: Objection. Lacks foundation, vague and ambiguous. 11:56:57 11:56:58

THE WITNESS: I believe that for purposes of brevity, that largely identical reports were -- were grouped together for visualization purposes. 11:57:03 11:57:06 11:57:31

BY MS. MOEHLMAN: 11:57:43

Q. And by grouping them together, would you consider that to be combining reports received? 11:57:43 11:57:46

MR. POLLACK: Objection. Vague and ambiguous. 11:57:52

THE WITNESS: Given a -- a plain meaning of the word "combining," sure. 11:57:57 11:58:01

EXHIBIT
F

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT G

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT H

Detecting Worms and Abnormal Activities with NetFlow, Part 1

Yiming Gong 2004-08-16

Editor's note: a French translation of this article (PDF) is also available.

Enterprise networks are facing ever-increasing security threats from worms, port scans, DDoS, and network misuse, and thus effective monitoring approaches to quickly detect these activities are greatly needed. Firewall and intrusion detection systems (IDS) are the most common ways to detect these activities, but additional technology such as NetFlow can be a valuable enhancement.

1. NetFlow overview

NetFlow is a traffic profile monitoring technology developed by Darren Kerr and Barry Bruins at Cisco Systems, back in 1996. As a de facto industry standard, NetFlow describes the method for a router to export statistics about the routed socket pairs, and it's now a built-in feature for most Cisco routers as well as Juniper, Extreme and some other vendor's routers and switches.

When a network administrator enables the NetFlow export on a router interface, traffic statistics of packets received on that interface will be counted as "flow" and stored into a dynamic flow cache.

1.1 What is "flow"?

Flow is defined as a unidirectional sequence of packets (which means there will be two flows for each connection session, one from the server to client, one from the client to server) between two endpoints. A flow can be identified by seven key fields: source IP address, destination IP address, source port number, destination port number, protocol type, type of services, and the router input interface. Any time after receiving a packet, a router will look for these seven fields and then make a decision: if the packet belongs to an existent flow, traffic statistics of the corresponding flow will be increased, otherwise a new flow entry will be created.

According to Cisco, as new flow is continuously created, the expired flow records will be exported by means of a UDP packet to a user-specified monitoring station if one of the following conditions occurs. The conditions are:

- The transport protocol indicates that the connection is completed (TCP FIN), and there is a small delay to allow for the completion of the FIN acknowledgment handshaking.
- Traffic inactivity exceeds 15 seconds.
- For flows that remain continuously active, flow cache entries expire every 30 minutes to ensure periodic reporting of active flows.

A number of network hardware vendors have implemented their version of NetFlow, but Version5 is now the most common. For a V5 datagram, every single UDP datagram contains one flow header and thirty flow records. Every flow record is made up of several fields, which include: the source and destination IP address, next hop address, input and output interface number, number of packet in the flow, total bytes in the flow, the source and destination port, the protocol, ToS, source and destination AS number, and TCP flags (Cumulative OR of TCP flags).

On the collection station, a flow file analyzer is needed to process the exported flow data in real time. It can be either commercial software/hardware or a station created with open source tools.

1.2 NetFlow versus intrusion detection systems

Looking through a flow record, you will find that there is no packet payload information in the flow field. This is one of the major differences with NetFlow as compared to a traditional IDS. A flow record doesn't contain any high-layer information, it just contains traffic profiles. As a result, this makes NetFlow lose the ability to dig deeply into packets and do any packet analysis work, yet there is still enough information to make some valuable conclusions from the data. The advantage to this approach is its high speed. Paying no attention to packet payloads

greatly reduces the processing overhead and makes NetFlow an extraordinarily good fit for busy, high-speed network environments. In addition, this characteristic makes NetFlow very useful in zero-day or "mutant attack" detection in cases where signature-based intrusion detection systems would fail.

Because flow data is coming directly from the router, a core element of any large network, NetFlow is capable of providing a unique view on the entire traffic of a network at the infrastructure level. It also proactive detection of network infrastructure security events.

If analyzed properly, NetFlow records will be very suitable for early worm and other abnormal network activity detection in large enterprise networks and service providers. In this paper, I will discuss some flow-based analysis methods on network security.

2. Flow-based analysis methods

2.1 Top N and Baseline

A baseline is a model describing what 'normal' network activity is according to some historical traffic pattern; all traffic that falls outside the scope of this established traffic pattern will be flagged as anomalous.

Trend and baseline analysis reports, commonly referred to as Top N and Baseline Analysis, is the most common and basic method of doing flow-based analysis. With this approach, attention is paid to flow records which have some "special high volume" characteristics, especially the value of those flow fields that deviate significantly from an established historical baseline.

Normally there are two ways to make use of Top N and Baseline methods: Top N sessions and Top N data.

2.1.1 Top N session

A Top N session means a single host produces an abnormally high volume of connection requests to a single destination or block of destinations, and the volume departs from the established baseline. The most likely reason for these activities are the presence of new worms, DoS/DDoS attacks, network scans or certain kinds of network abuses.

Normal clients connecting to the Internet should keep a relatively normal connection frequency to the outside. But if a host is infected with a worm, it will absolutely act different. It will always launch a large number of connection requests to the outside for its attempts to infect the next batch of victims, and as a result, the connection request numbers sent out will be significantly high.

For the same reason, when a lesser-skilled "script kiddie" is scanning a large block of addresses for certain vulnerable services, we will see especially high volume sessions sent out by that single IP address.

We can also use Top N session methods to detect many kinds of network abuses, such as checking the flow records for port 25 connection requests sent out by every single host in real time. In a given duration, for any host, if the statistics of port 25 requests are above a 'normal' value, it could be considered to be a spammer or someone infected with some kinds of email worm. It would be better for the Internet as a whole if service providers started using this technology and shut down the spammers upon detection.

2.1.2 Top N data

A second method of using Top N and Baseline methods is with Top N data. This can be defined as a consistantly large amount of network data transferred in a certain period of time between two network nodes or from a single node to a block of addresses.

The Top N hosts that transfer traffic data to or from the outside in an enterprise should be ranked into relatively fixed groups. If this pattern changes, and a new host suddenly appears in the Top N hosts matrix, an alert should be triggered.

Here is a example demonstrating Top N data methods that were used to track down a network security problem. One day, one of our customers reported a network bandwidth usage and congestion problem. We quickly enabled NetFlow on their upstream router's interface to collect egress traffic from their network, and had the flow data sent to our monitoring station. A few minutes later, a flow file was created. We analyzed the file with our flow-tools to generate a usage report for the top 20 hosts, sorted by octets. When the result displayed on console, we noticed that a host now siting in first place had abnormally high communication octets. A further examination of the flow records showed that the host sent out a huge number of requests to destination port 1434, so we now had the answer. The host was infected with the SQL slammer worm, and it almost ate up all their available bandwidth. After the customer patched the vulnerable machine, their network connection situation recovered.

2.2 Pattern Matching

Pattern matching is another method we can use to identify abnormal network activities when doing flow-based analysis. With this method, the flow records will be searched and those hosts associated with flow fields that seem "suspicious" based on our criteria will be flagged.

All the flow fields in a flow record can be used to do a pattern match, but the source and destination IP addresses, and the source and destination port numbers, are the most commonly used.

2.2.1 Port matching

Generally speaking, in order to launch an attack almost every attack should target a specific, functional port. For example, the SQL Slammer worm works on port 1434, the Netbus Trojan works on port 12345. Administrators can filter out all the flow records whereby the destination ports are equal to some specific ports, in order to find the corresponding attacks. This method is very easy to implement and can be used in most cases, although it may also produce false positives.

2.2.2 IP address matching

IP address matching is another method that can be used for security purposes with NetFlow analysis. There are several ways to make an IP address match, such as the following:

(A) Match IANA reserved addresses

The IANA has reserved large blocks of Internet address space which should not be used for global routing. If we find any flow record containing IANA reserved addresses, an alert should be triggered.

An important fact that the administrator must realize when performing IANA reserved address matches is that he can't trace back the potential host within the flow record if it is using spoofed IP addresses. At this point another flow field, Ifindex, should be used. We could check the corresponding router Ifindex number in the flow records to find the actual router interface where the flow comes from.

I've experienced an interesting case in which one of our customer's NetFlow records were appearing strange; the flow records showed a large number of connections whereby the source ports were all 80, the source addresses were 127.0.0.1, and the TCP flags of these flow records were all RST/ACK.

The following is an output example of flow-tools:

```
Sif  SrcIPaddress DIf  DstIPaddress     Pr SrcP DstP Pkts Octets StartTime          EndTime            Active  B/Pk  Ts  Fl
0059 127.0.0.1    005b 219.140.194.174  06 50   4f3   1    40     0721.21:58:00.593 0721.21:58:00.593 0.000   40    00  14
0059 127.0.0.1    005b 219.148.205.228  06 50   6ef   1    40     0721.21:57:56.533 0721.21:57:56.533 0.000   40    00  14
```

We can see that the source port (SrcP) is 50 in HEX, which equals 80 in decimal. And TCP flag (Fl) is 14 in HEX, and in the binary system it means 010100, which is TCP RST/ACK. Since the source IP address (SrcIPaddress) is a spoofed 127.0.0.1, where is the attacker coming from?

Using the router Ifindex (SIf) field in the flow records, the router interface where these packets came from was quickly identified. I informed the administrator who was in charge of the network on that interface, and after a little while he responded to me with the answer: a PC in his domain was broken in and had a DoS program installed. The program was designed to launch TCP port 80 DoS attacks with spoofed source IP addresses against a security website located in Guangdong, China, but the DNS A record of the website had been changed to 127.0.0.1. Thus, the attack packets were received by the PC itself, then reset to the spoofed source IP addresses.

(B). Match a special IP or IP list

There are always some default rules for any enterprise or ISP when performing flow-based abnormal detection. Some of those rules are based on:

- outbound traffic

  For an enterprise or ISP, any flow record where the IP source address is not part of their network domain for outbound traffic should be considered as abnormal.

- Inbound traffic

  For an enterprise or ISP, any flow record where the IP source addresses are part of their domain for inbound traffic should be considered abnormal.

- Fixed addresses

  Some kinds of abnormal activities may have one or more fixed IP addresses that contact is made with. For example, when the W32/Netsky.c worm spreads, it will send a DNS query to the following DNS servers,

  145.253.2.171, 151.189.13.35, 193.141.40.42, 193.189.244.205, 193.193.144.12, 193.193.158.10, 194.25.2.129, 194.25.2.129, 194.25.2.130, 194.25.2.131, 194.25.2.132, 194.25.2.133, 194.25.2.134, 195.185.185.195, 195.20.224.234, 212.185.252.136, 212.185.252.73, 212.185.253.70, 212.44.160.8, 212.7.128.162, 212.7.128.165, 213.191.74.19, 217.5.97.137, 62.155.255.16

  Therefore, any flow record in which the destination address is found to be in this list and the destination port is also UDP 53 should raise an alert, and future analysis is then needed.

3.0 Concluding part one

This concludes the first of our two-part series. Check back in two week's time where we'll continue the discussion of NetFlow. In part two, we'll look at how to filter our flow results via TCP flags, we'll discuss some ICMP issues, and then discuss some of the various tools that exist to help implement and analyze our NetFlow solution. Stay tuned.

About the author

Yiming Gong has worked for China Telecom for more than 5 years as a senior system administrator, and now he works as a Technical Manager in China Telecom System Integration Co.Ltd. He also has a personal homepage focusing on network/system security.

Comments on this article can be sent to the editor.

EXHIBIT
I

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT
J

RealSecure invalidates the indicated claims under 35 U.S.C. § 102(b)

"RealSecure" invalidity is shown both by the publications listed below and by the RealSecure product they identify and describe. RealSecure was on-sale more than a year before SRI's priority date.

- RealSecure 1.2.2: User Guide and Reference Manual (9/11/97) (ISS 312059-312076, ISS 312114-312128 and ISS 312134-312157)
- RealSecure 1.2: User Guide and Reference Manual (1997) (ISS 25469-25566)

  [Attached as Exhibit F to the 6/16/06 Declaration of Don Hall]
- RealSecure 1.1: User Guide and Reference Manual (3/97) (ISS 25387-25463)

  [Attached as Exhibit E to the 6/16/06 Declaration of Don Hall]
- RealSecure 1.0: User Guide and Reference Manual (1996) (ISS 354437-354465)

  [Attached as Exhibit D to the 6/16/06 Declaration of Don Hall]
- RealSecure Release 1.0 for Windows NT 4.0: A User's Guide and Reference Manual (ISS 02126117-02126244)

  [Attached as Exhibit C to the 6/16/06 Declaration of Don Hall]
- *More About RealSecure: General Description and Comparison to Existing Systems* (available at least as early as 07/21/1997, see archive.org) (ISS 357169-357178)

  [Attached as Exhibit H to the 6/16/06 Declaration of Don Hall]
- *Frequently-Asked Questions About RealSecure* (last updated 5/30/1997, available at least as early as 07/21/1997, see archive.org) (ISS 357179-357193)

  [Attached as Exhibit L to the 6/16/06 Declaration of Don Hall]
- *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security* (available at least as early as 1/20/98, see archive.org) (ISS 357242-357259)

  [ISS 357245-357259 is attached as Exhibit K to the 6/16/06 Declaration of Don Hall]

- *Frequently Asked Questions About Real-Secure* (last updated 10/21/97, available at least as early as 1/20/98, see archive.org) (ISS 357217-357227)
- RealSecure Press Releases (ISS 357164-357165, ISS 357262 and ISS 357263)

  [ISS 357164-357165 is attached as Exhibit J to the 6/16/06 Declaration of Don Hall]

  [ISS 357262 is attached as Exhibit B to the 6/16/06 Declaration of Don Hall]

  [ISS 357263 is attached as Exhibit A to the 6/16/06 Declaration of Don Hall]
- RealSecure Release Dates Table (ISS 358384)
- Articles relating to RealSecure (ISS_02125861- ISS _02125902).

  [ISS_02125872-875 is attached as Exhibit G to the 6/16/06 Declaration of Don Hall]

  [ISS_02125861-864 is attached as Exhibit N to the 6/16/06 Declaration of Don Hall]

  [ISS_02125900-902 is attached as Exhibit O to the 6/16/06 Declaration of Don Hall]

The text included herein are merely representative samples of the disclosure in the asserted reference. ISS reserves the right to supplement these disclosures.

Additional information is contained in:

- R. Power and R. Farrow, *Detecting Network Intruders*, Network Magazine, pp. 137-38, October 1997 (ISS 341748-341751)

'203 Patent

'203 Claim 1	RealSecure
1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:	RealSecure "is designed to be an enterprise-wide solution, monitoring the company's network at many different points of failure 24 hours a day, 7 days a week. RealSecure™ is composed of two parts: a filtering engine that watches and actively manages the network and GUI front-end that reports events and allows the user to configure the engine's scope. Multiple engines can be run on machines near critical points in the network, such as a firewall or a sensitive LAN. *More About RealSecure: General Description and Comparison to Existing Systems*, p. 1 (on ISS's website as of 07/21/1997). "RealSecure uses a distributed architecture. The RealSecure engine performs its filtering and monitoring functions on a given network segment. The RealSecure management console displays and logs the data and acts as a centralized engine management point." *Frequently-Asked Questions About RealSecure*, p. 5, Q. 10 (on the ISS website by 07/21/1997). "RealSecure, from Internet Security Systems, is a real-time monitoring, attack recognition, and response system. It monitors packet flow over a network in real time and analyzes packets for known attack patterns and unauthorized activity using a rule-based approach. When it detects an attack, it reacts automatically according to its configuration." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security*, p. 6 (available at least as early as 1/20/98, see archive.org). **"Advanced architecture** RealSecure consists of three components: - *Recognition engine.* This component monitors the network in real time, detecting and reporting attacks. It reports events to the Administrator's Module. - *Response Engine.* This component reacts automatically to recognized attack events, triggering prespecified actions ranging from logging attacks and alerting the

	administrator to terminating offending connections. - *Administrator's Module.* This component provides GUI (graphical user interface) management of the Recognition and Response engines. The Administrator's Module can monitor and manage all recognition and response engines from a single GUI, simplifying network management." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 7.
deploying a plurality of network monitors in the enterprise network;	"Q10: How is RealSecure deployed across the enterprise network? A: RealSecure uses a distributed architecture. The RealSecure engine performs its filtering and monitoring functions on a given network segment. The RealSecure management console displays and logs the data and acts as a centralized engine management point. Many RealSecure engines can report to a single management console. As engines detect unauthorized activity they take the appropriate action and then send a message to the management console so that the administrator can see what has happened. Engines can also upload their log files and databases to the management console periodically, so that the network administrator has a centralized report of network activity. With regard to placement of RealSecure engines, the best rule is to place a RealSecure engine on each segment where there is critical data to protect or a set of users that should be monitored. Note that a RealSecure engine will only see the traffic that is on the local network segment. Since routers prevent traffic from being copied to inappropriate segments, several RealSecure engines might be needed for complete coverage of network activity. The following figure shows a sample deployment of RealSecure. • There is a RealSecure engine behind the firewall monitoring the traffic flow on and off the network. This engine will detect unauthorized activity from the Internet. It will also help the administrator detect misconfigurations in the firewall. Many companies also place a RealSecure engine outside the firewall in the DMZ, to protect the external web server and to analyze external traffic. • There is a RealSecure engine on the Sales subnet because this is where the sales database resides. • The RealSecure engine on the Marketing subnet protects the business plans and marketing strategies on the systems in that department.

	• The Engineering subnet also supports a RealSecure engine because the source code archive is kept there. • All of these engines can report to a single management console. That console might be located on the company's internal network or at the headquarters across the internet or at a Network Operations Center staffed by a service organization." *FAQ About RealSecure*, pp. 5-6, Q10 (5/30/97). "As Figure 1 shows, an organization can place multiple Recognition engines in strategic locations in its network topology. An organization can also use multiple Recognition engines in parallel at a single location to accommodate high bandwith connections, such as T3 access to the Internet." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 8.
detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet};	"Q2: What kinds of network events does RealSecure recognize? A: RealSecure recognizes two types of network occurrences: Attacks Network activity patterns indicating that someone may be engaged in unauthorized or undesirable activity involving the systems and/or data on your network…" *FAQ About RealSecure*, p. 2-3, Q2 (5/30/97). "The administrator can configure the Recognition engine to implement specific security policies. The engine can react to (or ignore) connections based on specific packet types, source, and destination IP addresses or address ranges, port numbers, or particular types of attack patterns. This flexibility enables the administrator to set up custom monitoring for individual hosts and networks." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* pp. 7-8. *See also RealSecure 1.1: User Guide and Reference Manual* (March, 1997). This appendix summarizes RealSecure's features and attack signatures, some of which are quoted below. **Network Packet Data Transfer Commands**:

	"RealSecure can monitor many network services, including ... file transfer ..." *RealSecure 1.2.2: User Guide and Reference Manual,* Chapter 1, p. 1 (September 11, 1997). **"FTP GET File Decoding** Files being transferred from the source host to the destination host use a GET command in order to transfer the files. FTP GET decoding discovers all files that are being transferred to the source host over FTP." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 4. *See also RealSecure 1.1: User Guide and Reference Manual,* pp. A-10 to A-11 (March, 1997). **"FTP PUT File Decoding** Files being transferred from the source host to the destination host use a PUT command in order to transfer the files. FTP PUT decoding discovers all files that are being transferred to the destination host over FTP." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 4. *See also RealSecure 1.1: User Guide and Reference Manual,* p. A-11 (March, 1997). **"HTTP GET Decoding** Pages, images, and all other information that is viewed through a Web browser on the World Wide Web are transferred through HTTP using the GET command. HTTP GET decoding discovers all Web pages that are being transmitted unsecurely to a machine. This allows an administrator to track, log, and view Web traffic on the network." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 4.

	See also RealSecure 1.1: User Guide and Reference Manual, p. A-12 (March, 1997). **Network Packet Data Volume:** **"Ping Flooding** A Ping Flood is an attempt to saturate a network with packets in order to slow or stop legitimate traffic going through the network. A continuous series of ICMP Echo Requests are made to a target host on the network, which then responds with an ICMP Echo Reply. The continuing combination of requests and replies will slow the network and cause legitimate traffic to continue at a significantly reduced speed or, in extreme cases, to disconnect." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 2. **"Ping Flooding** This check determines if more than **PingFloodPackets** are received in **PingFloodDelta** seconds. The default setting is 50 packets in 3 seconds. If the network is on a slow connection like 14.4 PPP, consider making this setting more sensitive. Otherwise the default value should suffice." *RealSecure 1.1 – User Guide and Reference Manual*, pp. 3-15 to 3-16 (3/97) **"SYN Flood** A standard TCP connection is established by sending a SYN packet to the destination host.. If the destination is listening for a connection on the specified port, it will respond with a SYN/ACK packet. The initial sender then replies to the SYN/ACK with an ACK packet, and the connection is established. When the SYN/ACK is sent back to the source, a block of memory is allocated to hold information about the state of the connection that is currently being established. Until the final ACK is received or a timeout is reached, this block of memory sits unused waiting for more information to be received from the source host. By sending numerous SYN packets to a host, the destination will exhaust the portion of memory it has set aside to deal with opening connections, and legitimate connections will no longer be able to connect to the host.

	This situation can be detected by the flood of SYN packets without accompanying responses. It can be corrected by sending the destination RST packets that correspond to the initial SYNs. This results in the destination host freeing up that block of memory and making room for a new legitimate connection." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security*, p. 3. "RealSecure detects SYN Floods by monitoring the TCP connections that are established and by setting thresholds for the number of outstanding connections on a given machine at a given time." *Frequently-Asked Questions about RealSecure*, p. 13, Q27 (on the ISS website at least as early as 07/21/1997). "**SYN Flood Check** A SYN flood is a Denial of Service attack created by filling up the listen queue of a machine so that there is no room for legitimate users to establish a connection. If this attack is detected, and the 'Kill' action is set, **RealSecure** will implement a random drop algorithm that frees up an entry in the listen queue for a legitimate connection. In order to optimize the effectiveness of this algorithm, it is necessary to set the advanced parameter 'SYNFloodHighWaterMark.' This is the number of SYNs to allow to wait in each machine's queue for a response before the random drop algorithm is implemented. ..." *RealSecure 1.1 – User Guide and Reference Manual*, pp. 3-16 to 3-16 (3/97); *See also id.* at A-7. "**E-Mail Qmail Rcpt Denial of Service Vulnerability Check** (requires filter for TCP port 25) This check recognizes a Denial of Service attack against a Qmail mail server that is caused by repetitively RCPT commands to the server. An advanced parameter 'Email_Qmail_Rcpt_Threshold' can be configured to adjust the number of Rcpts that are legitimately allowed in a session before triggering this as an exploit. The default value for this parameter is 65535." *RealSecure 1.0 - User Guide and Reference Manual*, p. 89.

	Network Connection Requests: **"SYN Flood** A standard TCP connection is established by sending a SYN packet to the destination host. If the destination is listening for a connection on the specified port, it will respond with a SYN/ACK packet. The initial sender then replies to the SYN/ACK with an ACK packet, and the connection is established. When the SYN/ACK is sent back to the source, a block of memory is allocated to hold information about the state of the connection that is currently being established. Until the final ACK is received or a timeout is reached, this block of memory sits unused waiting for more information to be received from the source host. By sending numerous SYN packets to a host, the destination will exhaust the portion of memory it has set aside to deal with opening connections, and legitimate connections will no longer be able to connect to the host. This situation can be detected by the flood of SYN packets without accompanying responses. It can be corrected by sending the destination RST packets that correspond to the initial SYNs. This results in the destination host freeing up that block of memory and making room for a new legitimate connection." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 3. "RealSecure detects SYN Floods by monitoring the TCP connections that are established and by setting thresholds for the number of outstanding connections on a given machine at a given time." *Frequently-Asked Questions about RealSecure*, p. 13, Q27 (on the ISS website at least as early as 07/21/1997). **"SYN Flood Check** A SYN flood is a Denial of Service attack created by filling up the listen queue of a machine so that there is no room for legitimate users to establish a connection. If this attack is detected, and the 'Kill' action is set, **RealSecure** will implement a random drop algorithm that frees up an entry in the listen queue for a legitimate connection. In order to optimize the effectiveness of this algorithm, it is necessary to set the

	advanced parameter 'SYNFloodHighWaterMark.' This is the number of SYNs to allow to wait in each machine's queue for a response before the random drop algorithm is implemented. ..." *RealSecure 1.1 – User Guide and Reference Manual*, pp. 3-16 to 3-16 (3/97) "**ARP Check** If someone attempts to contact a host that is powered down, multiple address request packets are set with no response. **ArpMaxUnAcked** sets how many requests are sent to an unresponsive host, before triggering an alarm." *RealSecure 1.1 – User Guide and Reference Manual*, p. 3-16. "**FTP GET File Decoding** (requires filter for TCP port 21) Files being transferred from the destination host to the source host use a GET command in order to transfer the files. FTP GET decoding discovers all files that are being transferred to the source host over FTP." *RealSecure 1.0 - User Guide and Reference Manual*, p. 90. "**NetBIOS Session Grant Decode** (requires filter for TCP port 139) This decode recognizes when a NetBIOS session that has an outstanding session request has been granted the session. This indicates a session has been successfully established between the two machines." *RealSecure 1.0 - User Guide and Reference Manual*, p. 99. **Network Connection Denials:** "**NetBIOS Session Reject Decode** (requires filter for TCP port 139) This decode recognizes when a NetBIOS session that has an outstanding session

	request has been rejected. When available, a reason for the rejection will be provided. This indicates an attempted session has been denied between the two machines." *RealSecure 1.0 - User Guide and Reference Manual*, p. 99.
generating, by the monitors, reports of said suspicious activity; and	"As events occur on the network, the engines send messages to the GUI to indicate the event and the level." *More About RealSecure, General Description and Comparison to Existing Systems*, p. 3 (on web as of 07/21/97). "- *Recognition engine.* This component monitors the network in real time, detecting and reporting attacks. It reports events to the Administrator's Module." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security*, p. 7. "Data from the engines to the management console includes: -- Event messages – indications that something interesting has happened." *About Real-Secure*, p. 7, A11 (last updated October 21, 1997, on the ISS Website by January 20, 1998).
automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors.	"The gui displays [the event messages] according to priority." *More About RealSecure, General Description and Comparison to Existing Systems*, p. 3 (on web as of 07/21/97). "Data from the engines to the management console includes: -- Event messages – indications that something interesting has happened." *About Real-Secure*, p. 7, A11 (last updated October 21, 1997, on the ISS Website by January 20, 1998).

	"RealSecure's console is the central place where you review the captured suspicious network activity. As you see in Screen 4, the interface has five windows. In the left window, you can see a hierarchical view of the source address, the destination address, events, or actions taken on those events. This window's NT Explorer-style tree view provides an easy to drill down to the capture information. The three top windows on the right (High Priority, Medium Priority, and Low Priority) display each type of captured event according to its definable priority level." Screen produces a configurable tree-like display using a choice of source, destination, event type, and action (response taken). At each node of the tree, the number of reports satisfying the selected criteria is accumulated. The screen picture shows Events → Source → Destination → actions taken, with Event Type → counts The selected commonalities are Source, Destination, and Action. This accumulation of correlated date occurs on the console without any user activity. (RealSecure 1.0 for Windows NT, Windows NT Magazine, Oct. 1997, (2d page).
'203 Claim 2	
2. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.	*See RealSecure 1.1: User Guide and Reference Manual*, p. 4-3 (March, 1997) (showing *Network Events window*, grouping alerts by high, medium and low priorities). "RealSecure's console is the central place where you review the captured suspicious network activity. As you see in Screen 4, the interface has five windows. In the left window, you can see a hierarchical view of the source address, the destination address, events, or actions taken on those events. This window's NT Explorer-style tree view provides an easy to drill down to the capture information. The three top windows on the right (High Priority, Medium Priority, and Low Priority) display each type of

	captured event according to its definable priority level." Screen produces a configurable tree-like display using a choice of source, destination, event type, and action (response taken). At each node of the tree, the number of reports satisfying the selected criteria is accumulated. The screen picture shows Events → Source → Destination → actions taken, with Event Type → counts The selected commonalities are Source, Destination, and Action. This accumulation of correlated date occurs on the console without any user activity. (RealSecure 1.0 for Windows NT, Windows NT Magazine, Oct. 1997, (2d page). "With the reporting feature of RealSecure, "you can generate a report only on certain priorities of events, time periods and addresses." *RealSecure 1.1: User Guide and Reference Manual.* p. 5-4 (March, 1997). "RealSecure can generate meaningful reports from its event log files. These reports can include such information as the amount of data processed by a Web server each day, or the number of connections that were killed each day and from whom. The Administrator Module can display these reports in graphical form, such as bar or pie charts, for easy review and analysis. (See Figure 3)." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 11.
'203 Claim 3	
3. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.	"Q5: How does RealSecure respond to attacks? A: The actions taken upon detection of an attack or unauthorized activity are determined by the administrator. The administrator may choose from the following options: - Display a message indicating that the event occurred - View the session in real-time (or record for later playback)

	- Kill the connection automatically by sending a reset packet to each session participant - E-mail a notification to the administrator - Execute a user-specified program - Log the data related to the event for later reporting or playback." *Frequently-Asked Questions About RealSecure*, p. 4, (on the ISS website by 07/21/1997). "RealSecure's console is the central place where you review the captured suspicious network activity. As you see in Screen 4, the interface has five windows. In the left window, you can see a hierarchical view of the source address, the destination address, events, or actions taken on those events. This window's NT Explorer-style tree view provides an easy to drill down to the capture information. The three top windows on the right (High Priority, Medium Priority, and Low Priority) display each type of captured event according to its definable priority level." Screen produces a configurable tree-like display using a choice of source, destination, event type, and action (response taken). At each node of the tree, the number of reports satisfying the selected criteria is accumulated. The screen picture shows Events → Source → Destination → actions taken, with Event Type → counts The selected commonalities are Source, Destination, and Action. This accumulation of correlated date occurs on the console without any user activity. (RealSecure 1.0 for Windows NT, Windows NT Magazine, Oct. 1997, (2d page). *See also RealSecure 1.1 – User Guide and Reference Manual*, pp. 3-9 to 3-10 (3/97). "In addition to automatic response, the administrator can respond manually to reported attacks in a variety of ways using the Administrator's Module GUI: - *Request additional information.* The administrator can request the Engine to provide more detailed information on the reported attack. This information can include the

	packet source and packet data, such as e-mail headers. - *Instruct RealSecure to log the event.* If automatic event logging has not been configured for this event, the administrator can request RealSecure to log the reported event. - *Instruct RealSecure to kill the event.* If automatic session termination has not been configured for this event, the administrator can request RealSecure to terminate the session. The administrator can combine automatic and manual response to maintain the exact level of control required." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 11.
'203 Claim 4	
4. The method of claim 1, wherein the plurality of network monitors include an API for encapsulation of monitor functions and integration of third-party tools.	"Q29: Can RealSecure data be analyzed with a decision support system? A: Yes, if the decision support system is capable of reading an ODBC database." *Frequently-Asked Questions About RealSecure*, p. 14, Q. 29 (on the ISS website by 07/21/1997). Actions include "Run a user-specified program when the event occurs." RealSecure 1.1: User Guide and Reference Manual, ISS 25417, pp. 3-9. According to Mr. Joe Kleinwaechter, who I spoke with on April 18, customers implement calls to HP Openview from the engine with this mechanism.
'203 Claim 5	
5. The method of claim 1, wherein the enterprise network is a TCP/IP network.	"RealSecure can filter and monitor any TCP/IP protocol." *Frequently-Asked Questions about RealSecure*, p. 3 (Q. 3), on the ISS website as of 07/21/1997.
'203 Claim 6	

6. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.	"As Figure 1 shows, an organization can place multiple Recognition engines in strategic locations in its network topology. An organization can also use multiple Recognition engines in parallel at a single location to accommodate high bandwidth connections, such as T3 access to the Internet." *Real-Time Attack Recognition and Response: A Solution for Tightening Network Security,* p. 8. "To obtain maximum benefits, install RealSecure on a dedicated machine at an entry point to the network. Good places to consider would be at the Ethernet interface just inside the firewall or between the Internet router and the internal machines." *RealSecure 1.1 – User Guide and Reference Manual*, p. 2-1 (3/97). "[A] RealSecure engine will only see the traffic that is on the local network segment." *Frequently-Asked Questions About RealSecure*, p. 5, Q. 10 (on the ISS website by 07/21/1997).
'203 Claim 12	
12. An enterprise network monitoring system comprising:	See '203 claim 1
a plurality of network monitors deployed within an enterprise network, said plurality of network monitors detecting suspicious network activity based on analysis of network traffic data selected from the following categories: network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet};	See '203 claim 1
said network monitors generating reports of	See '203 claim 1

said suspicious activity; and	
one or more hierarchical monitors in the enterprise network, the hierarchical monitors adapted to automatically receive and integrate the reports of suspicious activity.	See '203 claim 1
'203 Claim 13	
13. The system of claim 12, wherein the integration comprises correlating intrusion reports reflecting underlying commonalities.	See '203 claim 2
'203 Claim 14	
14. The system of claim 12, wherein the integration further comprises invoking countermeasures to a suspected attack.	See '203 claim 3
'203 Claim 15	
15. The system of claim 12, wherein the plurality of network monitors include an application programming interface (API) for encapsulation of monitor functions and integration of third-party tools.	See '203 claim 4
'203 Claim 16	
16. The system of claim 12, wherein the enterprise network is a TCP/IP network.	See '203 claim 5
'203 Claim 17	
17. The system of claim 12, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.	See '203 claim 6

'615 Patent:
Invalidity If Claims Construed to Cover ISS Products
(Asserted Claims: 1-6 and 13-18)

RealSecure

U.S. Patent No. 6,711,615

'615 Claim 1	RealSecure	Comments
1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:	See '203 claim 1	
deploying a plurality of network monitors in the enterprise network;	See '203 claim 1	
detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};	See '203 claim 1 **Network Connection Acknowledgments**: **"IP Half Scan** A standard TCP connection is established by sending a SYN packet to the destination host. If the destination is waiting for a connection on the specified port, it will respond with a SYN/ACK packet. The initial sender then replies to the SYN/ACK with an ACK packet, and the connection is established. If the destination host is not waiting for a connection on the specified port, it will respond with an RST packet instead of a SYN/ACK. Most system logs do not log completed connections until the final ACK packet is received from the source. Sending an RST packet instead of the final ACK results in the connection never actually being established; so no logging takes	

	place. Because the source can identify whether the destination host sent a SYN/ACK or an RST, an attacker can determine exactly what ports are open for connections, without the destination ever being aware of probing." *See also RealSecure 1.1: User Guide and Reference Manual,* p. A-5 (March, 1997). **"ARP Check** If someone attempts to contact a host that is powered down, multiple address request packets are set with no response. **ArpMaxUnAcked** sets how many requests are sent to an unresponsive host, before triggering an alarm." *RealSecure 1.1 – User Guide and Reference Manual*, p. 3-16. **Network Packets Indicative of Well-Known Network Service Protocols**: **"IP Unknown Protocol** A standard IP packet contains an 8-bit protocol field. Common values for this field include 6 (TCP), 17 (UDP), and 1 (ICMP). Attackers sometimes use a non-standard value for this field, in order to exchange data between machines without logging mechanisms detecting the data that is being transmitted." *RealSecure 1.1 – User Guide and Reference Manual*, p. A-6 (3/97). **"HTTP Java Decoding** ... This decoding recognizes when a web browser attempts to obtain a file containing Java bytecode. this should only occur if a user has Java enabled on their web browser." *RealSecure 1.1: User Guide and Reference Manual,* p. A-13 (March, 1997).	
generating, by the monitors, reports of said suspicious activity; and	See '203 claim 1	
automatically receiving and integrating	See '203 claim 1	

the reports of suspicious activity, by one or more hierarchical monitors.		
'615 Claim 2		
2. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.	See '203 claim 2	
'615 Claim 3		
3. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.	See '203 claim 3	
'615 Claim 4		
4. The method of claim 1, wherein the plurality of network monitors include an API for encapsulation of monitor functions and integration of third-party tools.	See '203 claim 4	
'615 Claim 5		
5. The method of claim 1, wherein the enterprise network is a TCP/IP network.	See '203 claim 5	
'615 Claim 6		
6. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways routers, proxy servers}.	See '203 claim 6	
'615 Claim 13		
13. An enterprise network monitoring system comprising:	See '615 claim 1	

a plurality of network monitors deployed within an enterprise network, said plurality of network monitors detecting suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data volume, network connection requests, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};	See '615 claim 1	
said network monitors generating reports of said suspicious activity; and	See '615 claim 1	
one or more hierarchical monitors in the enterprise network, the hierarchical monitors adapted to automatically receive and integrate the reports of suspicious activity.	See '615 claim 1	
'615 Claim 14		
14. The system of claim 13, wherein the integration comprises correlating intrusion reports reflecting underlying commonalities.	See '615 claim 2	
'615 Claim 15		

15. The system of claim 13, wherein the integration further comprises invoking countermeasures to a suspected attack.	See '615 claim 3	
'615 Claim 16		
16. The system of claim 13, wherein the plurality of network monitors include an application programming interface (API) for encapsulation of monitor functions and integration of third-party tools.	See '615 claim 4	
'615 Claim 17		
17. The system of claim 13, wherein the enterprise network is a TCP/IP network.	See '615 claim 5	
'615 Claim 18		
18. The system of claim 13, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.	See '615 claim 6	